Document made available under the Patent Cooperation Treaty (PCT)

International application number: PCT/GB05/001332

International filing date: 06 April 2005 (06.04.2005)

Document type: Certified copy of priority document

Document details: Country/Office: US
Number: 10/818,685
Filing date: 06 April 2004 (06.04.2004)

Date of receipt at the International Bureau: 18 May 2005 (18.05.2005)

Remark: Priority document submitted or transmitted to the International Bureau in compliance with Rule 17.1(a) or (b)

PCT/GB 2005 / 0 0 1 3 3 2

PA 1306990

THE UNITED STATES OF AMERICA

TO ALL TO WHOM THESE PRESENTS SHALL COME:

UNITED STATES DEPARTMENT OF COMMERCE

United States Patent and Trademark Office

April 13, 2005

THIS IS TO CERTIFY THAT ANNEXED HERETO IS A TRUE COPY FROM THE RECORDS OF THE UNITED STATES PATENT AND TRADEMARK OFFICE OF THOSE PAPERS OF THE BELOW IDENTIFIED PATENT APPLICATION THAT MET THE REQUIREMENTS TO BE GRANTED A FILING DATE UNDER 35 USC 111.

APPLICATION NUMBER: *10/818,685*

FILING DATE: *April 06, 2004*

By Authority of the

COMMISSIONER OF PATENTS AND TRADEMARKS

M. Tarver

M. TARVER

Certifying Officer

22651 U.S. PTO

PTO/SB/05 (08-03)
Approved for use through 07/31/2006. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No.	920476-95843
First Inventor	Friskney
Title	Traffic Engineering in Frame-Based Carrier Networks
Express Mail Label	EV 402 895 299 US

19270 U.S. PTO
10/818685
040604

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

ADDRESS TO: Mail Stop Patent Application
Commissioner for Patents
P. O. Box 1450
Alexandria VA 22313-1450

1. [x] Fee Transmittal Form (e.g., PTO/SB/17) *(Submit an original, and a duplicate for fee processing)*
2. [ ] Applicant claims small entity status. See 37 CFR 1.27.
3. [x] Specification *(preferred arrangement set forth below)* [Total Pages 31]
   - Descriptive title of the invention
   - Cross Reference to Related Applications
   - Statement Regarding Fed sponsored R & D
   - Reference to sequence listing, a table, or a computer program listing appendix
   - Background of the Invention
   - Brief Summary of the Invention
   - Brief Description of the Drawings *(if filed)*
   - Detailed Description
   - Claim(s)
   - Abstract of the Disclosure
4. [x] Drawing(s) *(35 U.S.C. 113)* [Total Sheets 11]
5. Oath or Declaration [Total Pages 6]
   a. [x] Newly executed (original or copy)
   b. [ ] Copy from a prior application (37 CFR 1.63(d)) *(for continuation/divisional with Box 18 completed)*
      i. [ ] DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b).
6. [ ] Application Data Sheet. See 37 CFR 1.76
7. [ ] CD-ROM or CD-R in duplicate, large table or Computer Program *(Appendix)*
8. Nucleotide and/or Amino Acid Sequence Submission *(if applicable, all necessary)*
   a. [ ] Computer Readable Form (CRF)
   b. Specification Sequence Listing on:
      i. [ ] CD-ROM or CD-R (2 copies); or
      ii. [ ] Paper
   c. [ ] Statements verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

9. [x] Assignment Papers (cover sheet & document(s))
10. [ ] 37 CFR 3.73(b) Statement *(when there is an assignee)* [x] Power of Attorney
11. [ ] English Translation Document *(if applicable)*
12. [ ] Information Disclosure Statement (IDS)/PTO-1449 [ ] Copies of IDS Citations
13. [ ] Preliminary Amendment
14. [x] Return Receipt Postcard (MPEP 503) *(Should be specifically itemized)*
15. [ ] Certified Copy of Priority Document(s) *(if foreign priority is claimed)*
16. [ ] Nonpublication Request under 35 U.S.C. 122 (b)(2)(B)(i). Applicant must attach form PTO/SB/35 or its equivalent.
17. [ ] Other:

18. If a CONTINUING APPLICATION, *check appropriate box, and supply the requisite information below and in the first sentence of the specification following the title, or in an Application Data Sheet under 37 CFR 1.76:*

[ ] Continuation [ ] Divisional [ ] Continuation-in-part (CIP) of prior application No.: ______

Prior application information: Examiner ______ Group / Art Unit: ______

For CONTINUATION OR DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 5b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

19. CORRESPONDENCE ADDRESS

[x] *Customer Number:* 23644 *or* [ ] *Correspondence address below*

Name					
Address					
City		State		Zip Code	
Country		Telephone		Fax	

Name (Print/Type)	William M. Lee, Jr.	Registration No. (Attorney/Agent)	26,935
Signature	William M. Lee	Date	April 6, 2004

22651 U.S. PTO

PTO/SB/17 (10-03)
Approved for use through 07/31/2006. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

FEE TRANSMITTAL for FY 2004

Effective 10/01/2003. Patent fees are subject to annual revision.

☐ Applicant claims small entity status. See 37 CFR 1.27

TOTAL AMOUNT OF PAYMENT ($) $1,356.00

Complete if Known	
Application Number	
Filing Date	
First Named Inventor	Friskney
Examiner Name	
Art Unit	
Attorney Docket No.	920476-95843

METHOD OF PAYMENT *(check all that apply)*

☒ Check ☐ Credit card ☐ Money Order ☐ Other ☐ None

☐ Deposit Account:

Deposit Account Number: 12-0913

Deposit Account Name: Barnes & Thornburg

The Director is authorized to: *(check all that apply)*

☐ Charge fee(s) indicated below ☒ Credit any overpayments

☒ Charge any additonal fee(s) or any underpayment of fee(s)

☐ Charge fee(s) indicated below, except for the filing fee

to the above-identified deposit account.

FEE CALCULATION

1. BASIC FILING FEE

Large Entity Fee Code	Large Entity Fee ($)	Small Entity Fee Code	Small Entity Fee ($)	Fee Description	Fee Paid
1001	770	2001	385	Utility filing fee	770.00
1002	340	2002	170	Design filing fee	
1003	530	2003	265	Plant filing fee	
1004	770	2004	385	Reissue filing fee	
1005	160	2005	80	Provisional filing fee	

SUBTOTAL (1) ($) $770.00

2. EXTRA CLAIM FEES FOR UTILITY AND

			Extra Claims		Fee from below		Fee Paid
Total Claims	29	-20** =	9	X	18.00	=	162.00
Independent Claims	7	- 3** =	4	X	86.00	=	344.00
Multiple Dependent						=	

Large Entity Fee Code	Large Entity Fee ($)	Small Entity Fee Code	Small Entity Fee ($)	Fee Description
1202	18	2202	9	Claims in excess of 20
1201	86	2201	43	Independent claims in excess of 3
1203	290	2203	145	Multiple dependent claim, if not paid
1204	86	2204	43	** Reissue independent claims over original patent
1205	18	2205	9	** Reissue claims in excess of 20 and over original patent

SUBTOTAL (2) ($) $506.00

***or number previously paid, if greater; For Reissues, see above*

FEE CALCULATION (continued)

3. ADDITIONAL FEES

Large Entity Fee Code	Large Entity Fee ($)	Small Entity Fee Code	Small Entity Fee ($)	Fee Description	Fee Paid
1051	130	2051	65	Surcharge - late filing fee or oath	
1052	50	2052	25	Surcharge - late provisional filing fee or cover sheet	
1053	130	1053	130	Non - English specification	
1812	2,520	1812	2,520	For filing a request for *ex parte* reexamination	
1804	920*	1804	920*	Requesting publication of SIR prior to Examiner action	
1805	1,840*	1805	1,840*	Requesting publication of SIR after Examiner action	
1251	110	2251	55	Extension for reply within first month	
1252	420	2252	210	Extension for reply within second month	
1253	950	2253	475	Extension for reply within third month	
1254	1,480	2254	740	Extension for reply within fourth month	
1255	2,010	2255	1,005	Extension for reply within fifth month	
1401	330	2401	165	Notice of Appeal	
1402	330	2402	165	Filing a brief in support of an appeal	
1403	290	2403	145	Request for oral hearing	
1451	1,510	1451	1,510	Petition to institute a public use proceeding	
1452	110	2452	55	Petition to revive - unavoidable	
1453	1,330	2453	665	Petition to revive - unintentional	
1501	1,330	2501	665	Utility issue fee (or reissue)	
1502	480	2502	240	Design issue fee	
1503	640	2503	320	Plant issue fee	
1460	130	1460	130	Petitions to the Commissioner	
1807	50	1807	50	Processing fee under 37 CFR § 1.17(q)	
1806	180	1806	180	Submission of Information Disclosure Statement	
8021	40	8021	40	Recording each patent assignment per property (times number of properties)	80.00
1809	770	2809	385	Filing a submission after final rejection (37 CFR § 1.129(a))	
1810	770	2810	385	For each additional invention to be examined (37 CFR § 1.129(b))	
1801	770	2801	385	Request for Continued Examination (RCE)	
1802	900	1802	900	Request for expedited examination of a design application	

Other fee (specify) ____________

*Reduced by Basic Filing Fee Paid

SUBTOTAL (3) ($) $80.00

SUBMITTED BY

Complete *(if applicable)*

Name *(Print/Type)*	William M. Lee, Jr.	Registration No. (Attorney/Agent)	26,935	Telephone	(312) 214-4800
Signature	William M. Lee			Date	April 4, 2004

WARNING: Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

"Express Mail" mailing label number
EV 402 895 299 US
Date of deposit: April 6, 2004
I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. §1.10 on the date indicated above and is addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450

Jennifer J. Ramirez
(Typed or printed name of person mailing paper or fee)

(Signature of person mailing paper or fee)

16670IDUS01U: FRISKNEY, BRAGG, PARRY, ASHWOOD-SMITH, ALLAN

TRAFFIC ENGINEERING IN FRAME-BASED CARRIER NETWORKS

FIELD OF THE INVENTION

The present invention relates to methods of, software for and apparatus for enabling traffic engineering in carrier networks.

BACKGROUND TO THE INVENTION

For many years now, telecommunications carriers have been deploying packet-switched networks in place of circuit-switched networks for reasons of efficiency and economy. Packet-switched networks such as Internet Protocol (IP) or Ethernet networks are intrinsically connectionless in nature and as a result suffer from Quality of Service (QoS) problems. Customers value services which are guaranteed in terms of bandwidth and QoS.

Carriers may use Multi-Protocol Label Switching (MPLS) over a layer 2 network to create connection-oriented label switched paths (or tunnels) across the intrinsically connectionless network, and thereby to provide guaranteed QoS and bandwidth services to customers. However, MPLS is a relatively unstable standard and carriers desire an alternative.

It is desired to use Ethernet switches in carrier's networks. Use of Ethernet switches in carrier's networks would have the advantages of interoperability (mappings between Ethernet and other frame/packet/cell data structures such as IP and ATM are well known) and economy (Ethernet switches are relatively inexpensive compared to IP routers, for example).

However, the behaviour of conventional switched Ethernet networks is incompatible with carriers' requirements for providing guaranteed services to customers. Carriers need networks to be meshed for load balancing and resiliency – ie there must be multiple paths across it – and the ability to perform traffic engineering – ie the ability of the network operator to control the provision of explicitly routed variable bandwidth connections (or tunnels) through which traffic may be directed.

In contrast, conventional Ethernet networks must be simply-connected - ie there must be one and only one path between each and every node of the network. As a consequence, conventional Ethernet networks do not have support for network-wide load balancing, suffer from resiliency problems and cannot support traffic engineering.

Spanning tree protocols are known which enable a physically meshed Ethernet network to be logically transformed into a simply-connected network by detecting physical loops and logically disabling connections to break up the loops. Spanning tree protocols are also known which are able to detect failure of a physical connection (thereby partitioning the fully-connected network) and automatically restore one or more previously-disabled physical connections so as to re-connect the network. This provides a degree of resiliency. However, carriers are capable of and so desire to plan their network traffic routes to achieve much higher resiliency, flexibility and efficiency than spanning tree can achieve. This routing can most easily be achieved by segregating the traffic into connections whose routes are determined as part of this planning process.

Virtual Bridged LANs (or VLANs) are described in the Institute of Electrical and Electronics Engineers (IEEE) standard 802.1Q, 2003 Edition. Figure 1 shows a conventional VLAN 10 split up into a plurality of component LANs 12 and connected via VLAN-aware Media Access Control (MAC) bridges 14. Component LANs 12 are typically provided for different communities of interest, such as users sharing a common server or having common network

protocol requirements. Unique identifiers (VLAN tags or VLAN IDs) are used to identify each component LAN. Broadcast traffic is broadcast only within component LANs. This helps to overcome the scalability issues of the Ethernet by separating the whole VLAN 10 into smaller broadcast domains. VLAN tags are used to distinguish between traffic for different component LANs when forwarding traffic on shared links between MAC bridges.

The Internet Engineering Task Force (IETF) has published an Internet Draft referred to as draft-kawakami-mpls-lsp-vlan-00:txt. This document describes the use of VLAN tags for label switching across Ethernet networks in a manner similar to use of MPLS labels for label switching over MPLS networks - VLAN tags are used as labels to mark traffic at an ingress point of a label switched path (LSP) as belonging to a Layer 2 tunnel, and VLAN-aware Ethernet switches in the network act as a VLAN label switched routers. Connections are formed using one or more LSPs. Intermediate nodes along the connection may optionally swap the inbound label to a different outbound label.

One problem with the method proposed in draft-kawakami-mpls-lsp-vlan-00.txt is that a maximum of 4094 unique VLAN tags are definable in 802.1Q compliant equipment. This limits the flexibility in and increases the complexity of provisioning connections across the network. Another problem is that connections may not easily be re-routed once provisioned without creating transitory loops.

Another problem is that the Frame Check Sequence (FCS) in Ethernet frames is computed over both the payload and header portions of the frame. Thus, with the method proposed in draft-kawakami-mpls-lsp-vlan-00.txt, every time a VLAN tag (ie a label) is swapped at the ingress or egress point of a LSP, the FCS needs to be recomputed since the VLAN tag will have changed. This requires performing a computation function over the entire Ethernet frame. Moreover, during the interval from when the original FCS is removed

and the new FCS added, the frame is vulnerable to corruption without the protection of any FCS.

Another problem with the 'label-swapping' approach proposed in draft-kawakami-mpls-lsp-vlan-00.txt is that it requires a "chain of correctness" in that forwarding relies on each local label-forwarded link on the LSP being correct, whereas conventional Ethernet which uses globally unique address information to perform forwarding. More importantly, from a practical perspective, 'label-swapping' behaviour represents a significant change from conventional Ethernet switch functionality, and current telecommunications standards.

SUMMARY OF THE INVENTION

The present invention relates to enabling traffic engineering in frame-based networks such as Ethernet networks. The term traffic engineering is used broadly in the present document to refer to functions for maintaining the quality of service of the customers' connections while permitting the owner to operate their network efficiently. Examples of this are ensuring that no link is over-loaded, load-balancing the connections evenly across the network, re-equalizing the load on the network by re-routing some existing connections, establishing protection mechanisms, performing traffic restoration actions, and so on.

According to the present invention, connections are established in the carrier network by configuring, in one or more network nodes, mappings for forwarding data frames such as Ethernet frames. The mappings are from a combination of a) a destination address corresponding to a destination node of a connection, such as a MAC address, and b) an identifier, such as a VLAN tag. The mappings are to selected output ports of the one or more nodes. By using the combination of destination address AND identifier, the mappings enable data frames belonging to different connections to be forwarded differentially (ie forwarded on different output ports) despite the different connections potentially having the same destination node. This

enables flexibility in routing connections – eg the ability to perform traffic engineering. The reader should note that the term address is used in this document to denote any means of identifying a network node or an ingress or egress interface of a network node.

According to a first aspect of the present invention, there is provided a method of establishing connections in a frame-based network, the method comprising the step of configuring, in one or more nodes of the network, first mappings for use in forwarding data frames, the first mappings being from a combination of a first destination address corresponding to a first destination node of the network, and a first identifier, the first mappings being to a selected output port of, or to respective selected output ports of each of, the one or more nodes, thereby establishing at least part of a first connection through the one or more nodes to the first destination node.

Advantageously, the present invention enables connections to be established in a frame-based network in a highly flexible manner enabling network-wide traffic engineering. Furthermore, the specific problems inherent in the method proposed in draft-kawakami-mpls-lsp-vlan-00.txt (processing overhead and vulnerability of frames to corruption) are overcome since no label swapping is performed.

In one embodiment, the method of the present invention includes configuring, in at least one of the nodes, a second mapping for use in forwarding data frames, the second mapping being from a combination of: a second destination address corresponding to a second destination node of the network, and a second identifier, the second mapping being to a selected output port of the at least one node, thereby establishing at least part of a second connection through the at least one node to the second destination node, the selected output ports of the at least one node being different for the first and second mappings, thereby enabling, at the at least one node, differential forwarding of data frames associated with the first and second connections.

Thus, advantageously, two connections may be established which converge in route at an intermediate node and then diverge again, for example.

In one embodiment, the first and second destination addresses and the first and second destination nodes are the same. Thus, for example, two connections may be established which converge at an intermediate node and then diverge, despite having the same destination node. This enables greater flexibility in setting up connections.

In one embodiment, the first and second identifiers are the same. Thus, for example, two connections may be established which converge at an intermediate node and then diverge, despite using the same identifier. Thus, limitations on the number of values identifiers can take do not significantly reduce flexibility in traffic engineering.

Preferably, the network is an Ethernet network and the one or more nodes are Ethernet switches. Preferably, the identifier is a VLAN tag. Advantageously, this enables traffic engineered carrier networks to be deployed using conventional and relatively inexpensive VLAN-aware Ethernet switches, albeit configured in an entirely novel and inventive manner.

In one embodiment, the configuration is performed by a control plane of the network. Thus, carriers have direct control over the establishment of traffic engineering connections in the network. Preferably, the control plane is ASON/ASTN.

A frame-based communications network comprising one or more nodes arranged to perform the method of the first aspect of the present invention set out above is also provided.

Software arranged to perform the method of the first aspect of the present invention set out above is also provided.

According to a second aspect of the present invention, there is provided a connection controller for establishing connections in a frame-based network,

the connection controller comprising: a signal generator capable of generating a first signal for configuring, in a transport node of the network, a first mapping for use in forwarding data frames, the first mapping being from a combination of: a first destination address corresponding to a first destination node of the network, and a first identifier, the first mapping being to a selected output port of the transport node, the first signal thereby establishing at least part of a first connection through the transport node to the first destination node.

According to a third aspect of the present invention, there is provided a method of establishing a connection in a frame-based network, the method comprising the steps of: configuring forwarding information in a plurality of nodes of the network the forwarding information enabling the nodes to forward data frames in dependence on a combination of a destination address and an identifier of the data frames.

According to a fourth aspect of the present invention, there is provided a method of data traffic engineering in a frame-based network, the method comprising the following steps: establishing a first and second connections in the network passing through a common switching node of the network, configuring the switching node to forward data frames differently in dependence on differences in either a destination address or an identifier of the data frames, thereby enabling data traffic engineering.

According to a fifth aspect of the present invention, there is provided a method of establishing connections in a frame-based network, the method comprising the step of: configuring, in each of a first plurality of nodes of the network, a first forwarding mapping from a first combination of a destination address and an identifier to a selected output port of a respective node of the first plurality of nodes.

According to a sixth aspect of the present invention, there is provided a connection controller for establishing connections in a frame-based network, the connection controller being arranged to configure a first forwarding

mapping in a transport node, the first mapping being from a first combination of a destination address and an identifier to a first output port of the transport node.

According to a seventh aspect of the present invention, there is provided a method of forwarding data frames in a frame-based network, the method comprising the steps of: establishing a first connection in the network, the first connection being associated with a first combination of a destination address and an identifier, and forwarding data frames in the network in dependence on a combination of a destination address and an identifier of the data frames.

Further aspects of the present invention are set out in the appended claims. Further advantages of the present invention will be apparent from the following description.

In order to show how the invention may be carried into effect, embodiments of the invention are now described below by way of example only and with reference to the accompanying figures in which:

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows a conventional Virtual Bridged LAN;

Figure 2 shows an arrangement of Ethernet switches forming a carrier network according to the present invention;

Figure 3 shows a control plane/transport plane architecture for controlling the Ethernet carrier network of Figure 1 according to the present invention;

Figure 4 shows the carrier Ethernet network of Figure 1 arranged to provide connectivity between customer sites according to the present invention;

Figure 5 shows how nodes of the control plane interact with Ethernet switches of the transport plane to establish a connection across carrier network according to the present invention;

Figure 6 is a flow diagram showing the use of VLAN tag and destination address to differentiate forwarding of data traffic in different connections across the carrier network, according to the present invention;

Figure 7 shows an example of differential forwarding for two traffic flows having the same source and destination provider edge nodes but different VLAN tags according to the present invention;

Figure 8 shows an example of differential forwarding for two traffic flows having the same source provider edge nodes and VLAN tags but different destination provider edge nodes according to the present invention;

Figure 9 shows an example of converged routing for two traffic flows having the same destination provider edge node and VLAN tags but different source provider edge node according to the present invention;

Figure 10 shows a sparse mode of broadcast operation for customer VPNs provisioned across a carrier network, according to the present invention.

Figure 11 shows a dense mode of broadcast operation for customer VPNs provisioned across a carrier network, according to the present invention.

DETAILED DESCRIPTION OF INVENTION

Embodiments of the present invention are described below by way of example only. These examples represent the best ways of putting the invention into practice that are currently known to the Applicant although they are not the only ways in which this could be achieved.

To support guaranteed QoS to customers, what is required is:

1) an at least partially meshed carrier network;

2) the ability to establish explicitly routed connections across the carrier network between any two edge nodes (traffic engineering); and

3) the ability to enforce any bandwidth restrictions applied to the connections.

The present invention is primarily concerned with enabling requirements 1) and 2) above in frame-based networks such as Ethernet networks. Requirement 3) may be achieved using conventional mechanisms such as admission control at the ingress nodes of connections (trusted-edge policing).

Figure 2 shows an arrangement of Ethernet switches and communications links forming a carrier network according to the present invention. Carrier network cloud 20 comprises Ethernet switches 22a, 22b, 24a, 24b, 26 and 28. Ethernet switches 22a, 22b and 26 are located at the edges of carrier network 20, whereas Ethernet switches 24a, 24b, and 28 are located in the core network. Communications links (shown as straight lines in Figure 2) are provided between Ethernet switches 22a, 22b, 24a, 24b, 26 and 28. These communications links may be relatively long distance links over optical equipment such as SONET/SDH equipment with Ethernet interfaces using Generic Framing Procedure (GFP) (ITU-T Recommendation G.7041/Y.1303).

Note that core network switches 24a, 24b, and 28 are fully-meshed – ie there is a direct communications link connecting each core network switch 24a, 24b, and 28 to each other. Edge network switches 22a, 22b and 26 are not fully-meshed but have at least one direct communication link to communications link to a core network switch 26. The reader will appreciate that the particular network arrangement described is exemplary. In general, carrier networks may be implemented with virtually any number of Ethernet switches which, according to the present invention, may be connected in a fully-meshed or partially-meshed manner.

Figure 4 shows how a carrier Ethernet network may provide connectivity between customer sites according to the present invention. Three customers having respective pairs of geographically distant Ethernet switches (40a and 40b, 42a and 42b, and 44a and 44b) are shown connected to carrier network 20 via edge Ethernet switches 22a and 22b respectively. The communications links between edge switches 22a and 22b and customer switches 40a, 40b, 42a, 42b, 44a, and 44b may be dedicated links such as T1 leased lines or access links such as digital Subscriber Lines (DSLs).

Carrier edge switches 22a, 22b and 26 may be logically separated according to Nortel Network's Logical Provider Edge (LPE) architecture. Each carrier edge switch may thus be logically separated into a single Provider Edge- (PE-) Core and one or more PE-Edge functions. The PE-Edge is the ingress/egress point at which customer traffic enters or leaves the provider network - ie carrier network 20. The PE-Core encapsulates incoming Ethernet traffic from the customer using Media Access Control (MAC) in MAC encapsulation and forwards the encapsulated traffic across the carrier network. Similarly the PE-Core decapsulates (strips) outgoing Ethernet traffic and forwards the stripped traffic on to the customer via the appropriate PE-Edge. VLAN tags are used to provide customer separation at the PE-Core with each different customer site connected to each edge switch having a unique VLAN tag. Stacked VLAN (ie VLAN in VLAN encapsulation) may be used to protect any VLAN tags used by the customer traffic.

For example, customer switch 42a may send Ethernet traffic over communications link 46a to PE-Edge of edge switch 22a. PE-Core of edge switch 22a encapsulates each Ethernet frame in a further Ethernet frame using the MAC address of edge switch 22a as the source address and the MAC address of the appropriate egress point – in this case edge switch 22b – as the destination address. The encapsulated traffic is forwarded across a connection established over communications links 48 of carrier network 20 to edge switch 22b. Note that connections will typically be trunked in the sense that traffic from multiple customers will be routed through the same

connection. At the PE-Core of edge switch 22b, the original frames are stripped out and sent over communications link 46b via PE-Edge of edge switch 22b to customer switch 42b.

The reader will appreciate that in alternative embodiments of the present invention the PE-Edge may also be physically separated from the PE-Core and may reside at customer premises whereas the PE-Core resides at a central office or Point of Presence (PoP) of the carrier. The reader will also appreciate that other edge switches 26 will have connections to customer sites and that customers may have be provided with connectivity between two or more geographically distant sites over carrier network 20.

It will now be described how carrier network 20 is arranged to establish connections through which to forward encapsulated Ethernet traffic. A connection may be defined as an entity configured in a network which provides transport of data from a source node to one or more sink nodes.

As described above, carrier network 20 must be at least partially-meshed – ie there must be multiple paths between at least some, and preferably all, nodes of the network. Thus, as will be explained below, Ethernet MAC address auto learning functionality should preferably be at least partially deactivated.

On start-up (or on re-start), conventional switched Ethernet networks behave like a "classic" Ethernet Local Area Networks (LANs) in that every Ethernet frame is broadcast across the entire network. Thus, every switch, receiving an Ethernet frame on one port, broadcasts the frame out on every other port. The process repeats as the frame is received by other switches thus broadcasting the frame across the entire network.

MAC address auto-learning functionality is provided to improve efficiency in switched Ethernet networks. Ethernet frames have source and destination MAC addresses corresponding to their source and destination Ethernet switches. When an Ethernet frame sent out by a source switch is received by

an intermediate or destination Ethernet switch, the receiving switch observes the port on which the frame was received and the source address of the frame. It then builds up a forwarding table for use in future frame switching. The forwarding table maps destination address to output port and is built up using the source address of a received frame and the input port on which it was received. Over time, the network builds up forwarding state enabling efficient switching of Ethernet frames.

It can thus be seen that conventional switched Ethernet networks using auto-learning must be simply-connected - ie there must be one and only one path between each and every node of the network. If there were multiple paths between any two nodes, the input port on which a frame is received from a source node would not be a reliable indicator of the correct output port to forward future traffic destined for that node. Inconsistencies in forwarding tables on Ethernet switches would result in looping of frames. Moreover, if there exists any loop in an part of the network then any broadcast packet will be continuously duplicated in that loop and the duplicates forwarded all over the whole network, limited only by the link capacities concerned. This inevitably results in catastrophic failure of the network.

According to the present invention, instead of using auto learning to configure forwarding tables in Ethernet switches, forwarding tables are directly configured using a novel Ethernet control plane. Figure 3 shows a control plane/transport plane architecture for controlling the Ethernet carrier network of Figure 1. The ITU-T Automatically Switched Transport Network (ASTN), sometimes known as the Automatically Switched Optical Network (ASON), may be used. The general architectural specification of the ASTN is set out in ITU-T Recommendation G.8080.

Control plane 30 comprises a number of connection controllers 32a, 32b, 34a, 34b, 36 and 38 corresponding to each of Ethernet switches 22a, 22b, 24a, 24b, 26 and 28 of carrier network 20 (not all connection controllers are labelled in Figure 3, for clarity). Control Plane 30 may be conceptually

thought of as lying 'above' transport plane 32 which comprises the Ethernet switches 22a, 22b, 24a, 24b, 26 and 28 of carrier network 20. Connection controllers (CCs) 34 are logical agents each corresponding to a respective Ethernet switch (which represent cross connects in ASTN terminology) in transport plane 32. Each CC controls the switching of its respective switch using Connection Control Interface (CCI) signalling (shown as dotted lines in Figure 3). CCI signalling is used to directly configure the forwarding tables used by Ethernet switches 22a, 22b, 24a, 24b, 26 and 28 of carrier network 20. CCs may communicate between themselves using Network to Network Interface (NNI). Typically, CCs will exchange information regarding their operational state and the state, in particular the capacity, of their communications links using NNI signalling. Other control plane functions such as heartbeat, ping and circuit monitoring may be provided using the ITU-T standard-in-preparation currently referred to as Y.17ethOAM.

While CCs 32a, 32b, 34a, 34b, 36 and 38 are logically separate from Ethernet switches 22a, 22b, 24a, 24b, 26 and 28, the reader will understand that they may be implemented in the same physical nodes. Furthermore, NNI signalling may take place over the same communications links used for transporting user traffic.

Figure 5 shows how control plane 30 interacts with transport plane 32 to establish a point-to-point connection across carrier network 20. Typically, the connection will be bi-directional, although this can simply be considered as the combination of two uni-directional point to point connections. A request to establish a connection specifying a required bandwidth and an explicit route across carrier network 20 is generated for example by a supervisory network management node (not shown) or distributed network management or function. The explicit route will have been determined in accordance with a conventional routing protocol taking into account the topology of the carrier network, the operational state of network resources and the bandwidth requirements of existing and possible future connections. The route to be

taken by the exemplary connection shown in Figure 5 spans Ethernet switches 22a, 24a, 24b and 22b over communications links 48.

The request to establish a connection is first sent to CC 32a. On receipt of the request, CC 32a checks whether the communications link between switches 22a and 24a has sufficient capacity to support the required bandwidth. If so, it forwards a connection setup request message 50 to CC 34a specifying the required bandwidth and explicit route. CC 34a then checks whether the communications link between switches 24a and 24b has sufficient capacity to support the required bandwidth. The process continues until the connection setup message request 50 reaches CC 32b. Along the route, CCs may optionally reserve bandwidth of their respective switches and communication links so as to avoid race conditions where competing connections are setup over the same resources.

When connection setup request message 50 reaches CC 32b, if there is sufficient bandwidth along the entire path to support the required connection, then CC 32b sends a connection setup response message 52 back to CC 34b, CC 34a and finally to CC 32a. As the connection setup response message 52 traverses the CCs, each CC sends CCI signalling 54 to its respective switch to configure the forwarding tables of each switch, thereby to establish the forwarding state required to setup the connection.

It will be appreciated that the mechanism for establishing connections across carrier network 20 described above is merely exemplary and other well-known mechanisms may be used.

How forwarding tables of the Ethernet switches of carrier network 20 are used to support connections is a key aspect of the present invention and will now be described in detail.

Typically, there will be many thousands or tens of thousands of connections established across a carrier network at any time. These connections will share the physical resources of the carrier network – ie the switches and

communications links. Thus, each switch will typically have a large number of connections established through it at any point in time. However, each switch must be able to forward data traffic according to the explicit route requirements of the specific connection through which that traffic is being sent. A likely scenario is that the carrier network will need to establish multiple connections from the same source nodes, multiple connections to the same destination nodes and multiple connections both from the same source nodes and to the same destination nodes. However, for traffic engineering, even the latter connections may need to be established through different routes across the network. Furthermore, these routes may need to converge and diverge again within the carrier network. To support such route flexibility in connections, what is required is that each switch be able to differentiate between data traffic travelling in different connections and forward accordingly.

However, conventional switched Ethernet is incapable of this. As described above, conventional Ethernet switches forward traffic based solely on a forwarding table (established through auto learning) mapping destination address to output port. As a result, a conventional Ethernet switch will not be able to differentiate between data traffic having the same destination address, although it may be associated with multiple different connections.

According to the present invention, VLAN tags are used in a novel manner to enable differentiation of connections established across a carrier network and thereby to enable differential forwarding. Ethernet switches of carrier network 20 are VLAN-aware and use the combination of destination address and VLAN tag to forward data traffic. This may be achieved by each Ethernet switch storing separate forwarding tables for each VLAN tag configured, the VLAN tag acting as an mapping (or indexing) to the forwarding tables and each forwarding table mapping destination address to output port as normal. Thus the group of forwarding tables provide a mapping from a combination of destination address and VLAN tag to output port.

Thus, as shown in Figure 6, on receiving an Ethernet frame (step 60), the switch first selects a forwarding table based on the VLAN tag contained in the frame (step 62). Then, the switch selects an output port based on the destination address contained in the frame (step 64). Finally, the switch forwards the frame on the selected output port (step 66).

Note that this functionality is similar to that performed by the prior art VLAN-aware bridges described above with reference to Figure 1. However, the use of VLAN tags to enable the establishment and differentiation of connections across a carrier network is believed to be entirely novel and inventive.

According to the present invention, VLAN tags and entries in forwarding tables corresponding to connections to be established across the carrier network are directly configured into the appropriate Ethernet switches using the connection setup process described above. Data traffic is associated with a particular connection on entry into the carrier network (more specifically at the ingress PE-Core) by giving the frames a selected VLAN tag as well as destination address (ie the MAC address of the egress PE-Core). Note that encapsulation will already have taken place and thus the raw Ethernet frames received from the customer will not be altered in this process.

Figures 7 and 8 show how the use of a combination of VLAN tag and destination address enables differentiation between connections. Figure 9 shows how lack of differentiation in the combination of VLAN tag and destination address necessitates convergence between connections. Each of Figures 7 to 9 show connections across a carrier network comprising 4 provider edge Ethernet switches 71, 72, 73 and 74 (corresponding to PE1, PE2, PE3, PE4), further Ethernet switches in core 78 including core Ethernet switch 75, and communications links between the core and edge switches (reference numerals omitted for clarity).

In Figure 7, connections 76 and 77 have both the same source address (edge Ethernet switch 71 - PE1) and destination address (edge Ethernet switch 73 – PE3). However, the routes that connections 76 and 77 traverse are different.

In particular, it can be seen that at core Ethernet switch 75, connections 76 and 77 converge and then immediately diverge. Despite the common destination address, core Ethernet switch 75 is able to differentiate frames belonging to connection 76 from frames belonging to connection 77 (and to forward them accordingly) on the basis of their different VLAN tags. Thus, data traffic in connection 76 has the VLAN tag 2, for example, whereas data traffic in connection 77 has the VLAN tag 1.

In Figure 8, connections 80 and 82 have both the same source address (edge Ethernet switch 71 - PE1) and are given the same VLAN tag (in this case the VLAN tag is 1), but have different destination addresses (connection 80 has edge Ethernet switch 73 – PE3 while connection 82 has edge Ethernet switch 74 – PE4). Again, the routes that connections 80 and 82 traverse are different. In particular, it can be seen that at core Ethernet switch 75, connections 80 and 82 converge and then follow the same path before diverging towards their destination points. Despite the common VLAN tags, core Ethernet switch 75 is able to differentiate frames belonging to connection 76 from frames belonging to connection 77 (and to forward them accordingly) on the basis of their different destination addresses.

From Figures 8 and 9 it can be seen that, differentiation between Ethernet frames belonging to different connections is achieved according to the combination of destination address and VLAN tag. A difference in either may be used to achieve differential forwarding required for connections.

Figure 9 shows how lack of differentiation in the combination of VLAN tag and destination address necessitates convergence between connections. In Figure 9, connections 90 and 92 have the same destination address (edge Ethernet switch 73 - PE3), and are given the same VLAN tag (in this case the VLAN tag is 1), but have different source address (connection 90 has edge Ethernet switch 71 - PE1 while connection 92 has edge Ethernet switch 72 – PE2). Again, the routes that connections 90 and 92 traverse are different, but this is only because the data traffic is injected into the carrier network from

different ingress points – ie edge Ethernet switches 71 and 72. Once the routes converge at core Ethernet switch 75, they stay converged until their destination at edge Ethernet switch 73. This is because they have the same destination address and VLAN tag and there is no way of differentiating them on the basis of the combination of destination address and VLAN tag. However, it should be noted that one (or both) of the data traffic flows may be re-routed by simply provisioning a new connection with a different VLAN tag and then using that VLAN tag in the MAC header of the Ethernet frames at the ingress point. This re-routing of data flows is hitless since the new connection may be established contemporaneously with the old connection and new Ethernet frames directed into the new connection while earlier frames are still in transit over the old connection. Thus, re-routing may occur without requiring any re-ordering of frames at the destination.

So far, only the establishment of point-to-point (ie unicast) connections have been described. However, according to the present invention, point-to-multipoint or multipoint-to-multipoint connections may also be established across Ethernet networks as will now be briefly described. Conventional Ethernet switches are capable of multicast service. Typically this is achieved by configuring the forwarding table with more than one output port (but not necessarily all output ports) for a given multicast destination address. According to the present invention, for relatively small scale multicast operation, a point-to-multipoint connection may be configured as described above but using a combination of VLAN tag and multicast address mapping to more than one output port (but not necessarily all output ports) of selected Ethernet switches. However, this approach is only suitable for relatively small scale multicast operation.

According to the present invention, where a carrier network wishes to support a large number of point-to-multipoint or multipoint-to-multipoint connections, Resilient Packet Ring (RPR) are emulated over the Ethernet MAC addressed network using multiple unicast connections established as described above. The following description is given in the context of a virtual private network

(VPN) service, i.e. where there is a limited community of interest for each data frame. Two modes of operation are envisaged: a sparse mode for many customers with few sites, and a dense mode for few customers with many sites. The detailed mechanisms are described in one of the Applicants' co-pending US Patent Application Serial Number 10/698,833 (Nortel Networks Reference 15877RO) entitled Virtual Private Networks Within A Packet Network Having A Mesh Topology which document is incorporated herein by reference. The dense and sparse modes of operation will now be briefly described with reference to Figures 10 and 11.

Figure 10 shows a sparse mode of broadcast operation for many customers with few sites. Figure 10 shows a part of carrier network 20 comprising a part of fully-meshed core network 100, PE-Core edge Ethernet switches 104 a to d and PE-Edge edge Ethernet switches 102. Broadcast traffic 106a is received at PE-Core switch 104b from a customer site. Note that this traffic is broadcast within the context of a particular customer VPN, but is multicast within the context of the carrier network as a whole. The traffic is encapsulated and placed onto an RPR emulated by 4 uni-directional connections 108 a to d. The four connections are established as point-to-point connections as described above. The traffic is forwarded across each connection in turn until it reaches the start point again at PE-Core switch 104b. On receipt of an encapsulated frame, each endpoint of the four connections determines whether to process the frame for distribution to the customer via PE-Edge edge Ethernet switches 102 to which it is connected. This is done on the basis of broadcast destination addresses contained in the frames, and the VPN membership of customer sites attached to these Ethernet switches. Processing the frames involves decapsulating them and replicating them as required to one or more of PE-Edge edge Ethernet switches 102. It can be seen that no bandwidth need be dedicated to broadcast traffic in the sparse mode of operation since the four point-to-point connections may be trunked – ie they may be used to carry non-broadcast data and other customer's data, whether broadcast or not.

Figure 11 shows a dense mode of broadcast operation for few customers with many sites. Figure 11 shows a part of carrier network 20 comprising a part of fully-meshed core network 100, PE-Core edge Ethernet switches 104 a to d and PE-Edge edge Ethernet switches 102 as with Figure 10. Broadcast traffic 110a is received at PE-Core switch 104b from a customer site. Note, as above, that this traffic is broadcast within the context of a particular customer VPN, but is multicast within the context of the carrier network as a whole. The traffic is encapsulated and forwarded over a uni-directional connection 110b to a core switch 116a. Uni-directional connection 110b may be trunked. At core switch 116a, the traffic is forwarded in over a bi-directional RPR 112 emulated by connections between core switches 116 a to d using a bidirectional connection between each pair of adjacent nodes. The RPR is dedicated to a particular customer's broadcast traffic and is not trunked. This is achieved by using a unique VLAN tag for forwarding in the RPR.

The traffic is forwarded around RPR 112 to each of the core switches 116 a to d in one direction or the other, whichever is shortest for each respective core switch. Each core switch broadcasts the received frames over uni-directional connections 114 a so that each of PE-Core switches 104 a to d receives the traffic. Then, as with the sparse mode of broadcast operation described above, each PE-Core switch determines whether to process the frame for distribution to the customer via PE-Edge edge Ethernet switches 102 to which it is connected. This is done on the basis of broadcast destination addresses contained in the frames and involves decapsulating and replicating them as required to one or more of PE-Edge switches 102 for onward transmission to the customer sites.

It has been described above how connections may be established over a meshed Ethernet carrier network through configuring forwarding tables in network nodes and how data may be forwarded over those connections. The reader will appreciate that connections may be removed by deleting the

configuration data from every node over which the connection was established. It is important that all such configuration data is removed to avoid network failure or inefficiency. As described above, the default behaviour of Ethernet switches on receiving a frame addressed to an unknown destination (ie where there is no forwarding state configured for that destination address) is to broadcast the frame out on all output ports. In simply–connected networks this behaviour is appropriate. However, with a meshed topology, this behaviour can be catastrophic. Through partial removal of connections (in particular where configuration data is left at ingress points of a connection but deleted at points further along the connections towards or including the egress point), it remains possible that Ethernet frames for the PE may enter the network but arrive at a point where there is no configuration data for forwarding them, resulting in undesirable broadcast behaviour. Furthermore, partial removal of connections may leave forwarding loops configured by accident.

One solution to the problem of partial removal of connections is to alter the behaviour of the Ethernet switches forming the carrier network so that instead of broadcasting unknown traffic, they discard packets and possibly issue an alarm, log or count the discarded packets. However, altering the basic behaviour of Ethernet switches may require a hardware modification. While possible, this is not preferable. However, conventional Ethernet switches generally provide a software configurable function called rate limitation. Preferably, at all or most switches of the carrier rate limitation is used to set a rate of zero, or a low rate if zero is not possible, for broadcast traffic including broadcast-on-unknown traffic.

Where this is not possible, other pre-emptive approaches to minimising the problems of partial removal of connections may be used. One approach is to use block lists. Conventional Ethernet switches provide a block list (typically of limited length) which may be used to specify certain destination MAC addresses such that received Ethernet frames addressed to these blocked

address will be discarded without forwarding. By blocking, at all or most nodes of the network, the MAC addresses of many (but not all) MAC addresses of provider edge nodes it is possible to minimise the potential dangers of partial removal of connections without over restricting the carrier's flexibility in establishing connections across the network. Notably, it is necessary to block different MAC address at different nodes of the network. Typically, at a given node, the block list will include only the MAC address for provider edge nodes to which no connections are likely to be established through that node. This approach is not easily scaleable with large networks (the limited number of entries in block lists may be exhausted by large numbers of provider edge nodes). However, note that to prevent loops it is only necessary to block rogue frames at one node in any loop. Thus, it is possible to "spread" the blocked destination addresses more thinly across the network and still provide a degree of protection from loops thereby making more efficient use of the limited capacity of block lists.

While it is the use of VLAN tags in the present invention that enables flexibility in establishing connections across the network, the failure to remove VLAN state fully leaves the potential for looping of traffic. In particular, the problem will arise where a logical loop is left configured for any single given VLAN tag – ie the output ports of nodes defining a physical loop are left configured with membership of any single VLAN. Thus, another pre-emptive approach to minimising the problems of partial removal of connections is to allocate connections to or from neighbouring or nearby provider edge nodes using mutually exclusive VLAN tag pools. Thus, for example all connections to or from provider edge node PE1 will be guaranteed to have a different VLAN tag to those to or from neighbouring provider edge node PE2. In this way, VLAN loops including both PE1 and PE2 cannot accidentally be formed through the partial removal of connections since by definition any state left configured in PE1 and PE2 will use different VLAN tags. This approach may be generalised by allocating connections to or from n adjacent provider edge nodes using n mutually exclusive VLAN tag pools. n is chosen to be

sufficiently large to segregate use of VLAN tag pools as much as possible while providing sufficient flexibility in connection establishment to or from any particular provider edge node (remembering that there are only 4094 possible VLAN tags). With smaller carrier networks it may be possible for each provider edge node to use a different VLAN tag pool. However, with larger carrier networks it will be necessary to re-use VLAN tag pools at topologically distant provider edge nodes otherwise flexibility in connection establishment will be compromised though VLAN tag pools being too small.

It will be appreciated that combinations of the above approaches to minimising the problems of partial removal of connections may be employed.

While embodiments have been described above with reference to the use of VLAN tags for enabling flexibility in establishing and differential forwarding of data traffic associated with different connections, the reader will appreciate that other tags or identifiers may be used.

Also, while embodiments have been described above with reference Ethernet networks and Ethernet frames, to the reader will appreciate that the present invention applies in general to any frame-based, packet-based or cell-based switching network whether at OSI layer 2 or layer 3 network. And to data structures including frames, packets and cells. In the following claims, the term frame-based network, or cognate terms, shall denote any such switching network and the term frame, or cognate terms, shall denote any such data structure.

The reader will appreciate that the methods described above may be implemented in the form of hardware or software operating on conventional data processing hardware.

CLAIMS

1. A method of establishing connections in a frame-based network, the method comprising the step of:

configuring, in one or more nodes of the network, first mappings for use in forwarding data frames, the first mappings being from a combination of:

a) a first destination address corresponding to a first destination node of the network, and

b) a first identifier,

the first mappings being to a selected output port of, or to respective selected output ports of each of, the one or more nodes, thereby establishing at least part of a first connection through the one or more nodes to the first destination node.

2. A method according to claim 1 comprising the further step of:

configuring, in at least one of the nodes, a second mapping for use in forwarding data frames, the second mapping being from a combination of:

a) a second destination address corresponding to a second destination node of the network, and

b) a second identifier,

the second mapping being to a selected output port of the at least one node, thereby establishing at least part of a second connection through the at least one node to the second destination node,

the selected output ports of the at least one node being different for the first and second mappings, thereby enabling, at the at least one node,

differential forwarding of data frames associated with the first and second connections.

3. A method according to claim 2, wherein the first and second destination addresses and the first and second destination nodes are the same.

4. A method according to claim 2, wherein the first and second identifiers are the same.

5. A method according to claim 2, wherein the at least one node associates data frames with one of the first and second connections in dependence on a destination address and identifier of the frame.

6. A method according to claim 1, wherein the network is an Ethernet network and the one or more nodes are Ethernet switches.

7. A method according to claim 6, wherein the identifier is a VLAN tag.

8. A method according to claim 1, wherein the configuration is performed by a control plane of the network.

9. A method according to claim 8, wherein the control plane is ASON/ASTN.

10. A method according to claim 1, wherein the network is at least partially meshed.

11. A frame-based communications network comprising one or more nodes arranged to perform the method of claim 1.

12. Software arranged to perform the method of claim 1.

13. A connection controller for establishing connections in a frame-based network, the connection controller comprising:

a signal generator capable of generating a first signal for configuring, in a transport node of the network, a first mapping for use in forwarding data frames, the first mapping being from a combination of:

a) a first destination address corresponding to a first destination node of the network, and

b) a first identifier,

the first mapping being to a selected output port of the transport node, the first signal thereby establishing at least part of a first connection through the transport node to the first destination node.

14. A connection controller according to claim 13, wherein the signal generator is capable of generating a second signal for configuring, in the transport node of the network, a second mapping for use in forwarding data frames, the second mapping being from a combination of:

a) a second destination address corresponding to a second destination node of the network, and

b) a second identifier,

the second mapping being to a selected output port of the transport node, the second signal thereby establishing at least part of a second connection through the transport node to the second destination node,

the selected output ports of the transport node being different for the first and second mappings, the first and second signals thereby enabling, at the transport node, differential forwarding of data frames associated with the first and second connections.

15. A connection controller according to claim 14, wherein the first and second destination addresses and the first and second destination nodes are the same.

16. A connection controller according to claim 14, wherein the first and second identifiers are the same.

17. A connection controller according to claim 13, wherein the network is an Ethernet network and the transport node is an Ethernet switch.

18. A connection controller according to claim 17, wherein the identifier is a VLAN tag.

19. A connection controller according to claim 13, wherein the network is at least partially meshed.

20. A connection controller according to claim 13 physically co-located with the transport node.

21. A frame-based network comprising the connection controller of claim 13.

22. A method of establishing a connection in a frame-based network, the method comprising the steps of:

configuring forwarding information in a plurality of nodes of the network the forwarding information enabling the nodes to forward data frames in dependence on a combination of a destination address and an identifier of the data frames.

23. A method of data traffic engineering in a frame-based network, the method comprising the following steps:

establishing a first and second connections in the network passing through a common switching node of the network,

configuring the switching node to forward data frames differently in dependence on differences in either a destination address or an identifier of the data frames, thereby enabling data traffic engineering.

24. A method of establishing connections in a frame-based network, the method comprising the step of:

configuring, in each of a first plurality of nodes of the network, a first forwarding mapping from a first combination of a destination address and an identifier to a selected output port of a respective node of the first plurality of nodes.

25. A method according to claim 24 comprising the further step of:

configuring, in at least one of the first plurality of nodes, a second forwarding mapping from a second combination of a destination address and an identifier to a selected output port of the at least one node,

wherein the first and second mappings at the at least one node respectively map to different output ports of the at least one node.

26. A connection controller for establishing connections in a frame-based network, the connection controller being arranged to configure a first forwarding mapping in a transport node, the first mapping being from a first combination of a destination address and an identifier to a first output port of the transport node.

27. A connection controller according to claim 26, the connection controller being arranged to configure a second forwarding mapping in the transport node, the second mapping being from a second, different combination of a destination address and an identifier to a second, different output port of the transport node, thereby to enable differential forwarding of data frames by the transport node in dependence on a combination of a destination address and an identifier of the data frames.

28. A method of forwarding data frames in a frame-based network, the method comprising the steps of:

establishing a first connection in the network, the first connection being associated with a first combination of a destination address and an identifier, and

forwarding data frames in the network in dependence on a combination of a destination address and an identifier of the data frames.

29. A method according to claim 28, wherein the first connection is routed through a switching node having different first and second output ports connected to other nodes of the network, the method comprising the further steps of:

establishing a second connection in the network through the switching node, the second connection being associated with a second combination of a destination address and an identifier, different to the first combination, and

the switching node forwarding data frames having the first combination of a destination address and an identifier on the first output port and forwarding data frames having the second combination of a destination address and an identifier on the second output port.

ABSTRACT

Include Figure 7

The invention relates to enabling traffic engineering in frame-based networks such as Ethernet networks. There is described a method of and connection controller for establishing connections (76, 77) in a frame-based communications network comprising nodes (71-75 and 78) such as Ethernet switches. The connections are established by configuring, in various of the nodes, mappings for forwarding data frames, such as Ethernet frames. The mappings are from a combination of a) a destination address corresponding to a destination node (73) of the connection and b) an identifier, such as a VLAN tag. The mappings are to selected output ports of the various nodes. By using the combination of destination address AND identifier, the mappings enable data frames belonging to different connections (76, 77) to be forwarded differentially (ie forwarded on different output ports) at a node (75) despite the different connections having the same destination node. This enables flexibility in routing connections – ie the ability to perform traffic engineering.

PATENT/DOCKET NO. 920476-95843
CUSTOMER NO. 23644

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re Application of:	§	Attorney Docket No. ________
Serial No.: n/a	§	
Filed: Herewith	§	Group Art Unit: n/a
Title: Traffic Engineering in Frame-Based Carrier Networks	§	Examiner: n/a

Commissioner for Patents
Washington, D.C. 20231

POWER OF ATTORNEY FOR
PATENT APPLICATION

Nortel Networks Limited, a Canadian corporation having principal offices at:

Nortel Networks Limited
2351 Boulevard Alfred-Nobel
St. Laurent, Quebec H4S 2A9, Canada

the owner by assignment of the entire right, title and interest to the invention for "Traffic Engineering in Frame-Based Carrier Networks", by Robert Friskney, Nigel Bragg, Simon Parry and Peter Ashwood-Smith, and in and to the application for patent and any Letters Patent, whether domestic or foreign, that may issue thereon, by virtue of the assignment (check as applicable):

- ☒ Recorded Concurrently Herewith
- ☐ Recorded on ____________
- ☐ Recorded at Reel _______ Frame _______
- ☐ Copy Attached

hereby appoints the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Thomas E. Smith (Reg. No. 18,243, Dennis M. McWilliams (Reg. No. 25,195), James R. Sweeney (Reg. No. 18,721), William M. Lee, Jr. (Reg. No. 26,935), Glenn W. Ohlson, (Reg. No. 28,455), David C. Brezina (Reg. No. 34,128), Jeffrey R. Gray (Reg. No. 33,391), Timothy J. Engling (Reg. No. 39,970), Gerald S. Geren (Reg. No. 24,528), Peter J. Shakula (Reg. No. 40,808), Robert F.I. Conte (Reg. No. 20,354), Howard B. Rockman (Reg. No. 22,190), Mark J. Nahnsen (Reg. No. 51,093), Mark A. Hagedorn (Reg. No. 44,731);

its attorneys with full power of substitution and revocation, to prosecute all domestic and foreign patent applications, including PCT and EPO filings, relating to said invention and to transact all business connected therewith, including signing of all papers on its behalf and making alterations and amendments.

The undersigned is the representative for the Assignee of the entire right, title, and interest in the patent application identified above, and is authorized to act on behalf of the Assignee.

NORTEL NETWORKS LIMITED

6 April 2004
Date

Name: **Daniel Hermele**

Patent Associate
Title

PTO/SB/01 (08-03)
Approved for use through 07/31/2006. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

DECLARATION FOR UTILITY OR DESIGN PATENT APPLICATION (37 CFR 1.63)

☑ Declaration Submitted With Initial Filing **OR** ☐ Declaration Submitted after Initial Filing (surcharge (37 CFR 1.16 (e)) required)

Attorney Docket Number	926476-45843
First Named Inventor	R FRISKNEY
COMPLETE IF KNOWN	
Application Number	
Filing Date	
Art Unit	
Examiner Name	

I hereby declare that:

Each inventor's residence, mailing address, and citizenship are as stated below next to their name.

I believe the inventor(s) named below to be the original and first inventor(s) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

TRAFFIC ENGINEERING IN FRAME-BASED CARRIER NETWORKS

(Title of the Invention)

the specification of which

☑ is attached hereto

OR

☐ was filed on (MM/DD/YYYY) ______ as United States Application Number or PCT International Application Number ______ and was amended on (MM/DD/YYYY) ______ (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment specifically referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in 37 CFR 1.56, including for continuation-in-part applications, material information which became available between the filing date of the prior application and the national or PCT international filing date of the continuation-in-part application.

I hereby claim foreign priority benefits under 35 U.S.C. 119(a)-(d) or (f), or 365(b) of any foreign application(s) for patent, inventor's or plant breeder's rights certificate(s), or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent, inventor's or plant breeder's rights certificate(s), or any PCT international application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application Number(s)	Country	Foreign Filing Date (MM/DD/YYYY)	Priority Not Claimed	Certified Copy Attached? Yes	No
			☐	☐	☐
			☐	☐	☐
			☐	☐	☐
			☐	☐	☐

☐ Additional foreign application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto.

PTO/SB/01 (06-03)
Approved for use through 07/31/2003. OMB 0851-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

DECLARATION — Utility or Design Patent Application

Direct all correspondence to: ☑ Customer Number: 23644	OR ☐ Correspondence address below		
Name BARNES & THORNBURG			
Address PO BOX 2786			
City CHICAGO	State ILLINOIS	ZIP 60690-2786	
Country UNITED STATES OF AMERICA	Telephone (312) 357 1313	Fax (312) 759 5646	

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. 1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

NAME OF SOLE OR FIRST INVENTOR:	☐ A petition has been filed for this unsigned inventor		
Given Name (first and middle [if any]) ROBERT	Family Name or Surname FRISKNEY		
Inventor's Signature	Date 5th April 2004		
Residence: City HARLOW	State ESSEX	Country UNITED KINGDOM	Citizenship UNITED KINGDOM
Mailing Address 78 ALBERT GARDENS			
City HARLOW	State ESSEX	ZIP CM17 9QG	Country UNITED KINGDOM
NAME OF SECOND INVENTOR:	☐ A petition has been filed for this unsigned inventor		
Given Name (first and middle [if any]) NIGEL	Family Name or Surname BRAGG		
Inventor's Signature Nigel L Bragg	Date 5 April 2004		
Residence: City WESTON COLVILLE	~~State~~ CAMBRIDGESHIRE	Country UNITED KINGDOM	Citizenship UNITED KINGDOM
Mailing Address HOMEWARDS CHAPEL ROAD			
City WESTON COLVILLE	State CAMBRIDGESHIRE	ZIP CB1 5NX	Country UNITED KINGDOM

☑ Additional inventors or a legal representative are being named on the 1 supplemental sheet(s) PTO/SB/02A or 02LR attached hereto.

PTO/SB/02A (08-03)
Approved for use through 08/31/2003. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

DECLARATION

ADDITIONAL INVENTOR(S)
Supplemental Sheet

Name of Additional Joint Inventor, if any: ☐ A petition has been filed for this unsigned inventor

Given Name (first and middle (if any))	Family Name or Surname		
SIMON	PARRY		
Inventor's Signature: Simon Parry			Date 5 APRIL 2004
HARLOW Residence: City	ESSEX State	UNITED KINGDOM Country	UNITED KINGDOM Citizenship
2 KINGSTON COTTAGES, MATCHING GREEN Mailing Address			
2 KINGSTON COTTAGES, MATCHING GREEN Mailing Address			
HARLOW City	ESSEX State	CM17 0PS Zip	UNITED KINGDOM Country

Name of Additional Joint Inventor, if any: ☐ A petition has been filed for this unsigned inventor

Given Name (first and middle (if any))	Family Name or Surname		
PETER	ASHWOOD-SMITH		
Inventor's Signature	Date		
HULL Residence: City	QUEBEC State	CANADA Country	CANADA Citizenship
20 DES GENEVRIERS Mailing Address			
20 DES GENEVRIERS Mailing Address			
HULL City	QUEBEC State	J9A 2V8 Zip	CANADA Country

Name of Additional Joint Inventor, if any: ☐ A petition has been filed for this unsigned inventor

Given Name (first and middle (if any))	Family Name or Surname		
DAVID	ALLAN		
Inventor's Signature	Date		
OTTAWA Residence: City	ONTARIO State	UNITED KINGDOM Country	CANADA Citizenship
852 FOREST ST. Mailing Address			
852 FOREST ST. Mailing Address			
OTTAWA City	ONTARIO State	K2B 5P9 Zip	CANADA Country

PTO/SB/01 (08-03)
Approved for use through 07/31/2006. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

DECLARATION FOR UTILITY OR DESIGN PATENT APPLICATION (37 CFR 1.63)

Attorney Docket Number	970476-95843
First Named Inventor	R FRISKNEY
COMPLETE IF KNOWN	
Application Number	
Filing Date	
Art Unit	
Examiner Name	

☑ Declaration Submitted With Initial Filing OR ☐ Declaration Submitted after Initial Filing (surcharge (37 CFR 1.16 (e)) required)

I hereby declare that:

Each inventor's residence, mailing address, and citizenship are as stated below next to their name.

I believe the inventor(s) named below to be the original and first inventor(s) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

TRAFFIC ENGINEERING IN FRAME-BASED CARRIER NETWORKS

(Title of the Invention)

the specification of which

☑ is attached hereto

OR

☐ was filed on (MM/DD/YYYY) ______ as United States Application Number or PCT International Application Number ______ and was amended on (MM/DD/YYYY) ______ (if applicable).

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment specifically referred to above.

I acknowledge the duty to disclose information which is material to patentability as defined in 37 CFR 1.56, including for continuation-in-part applications, material information which became available between the filing date of the prior application and the national or PCT international filing date of the continuation-in-part application.

I hereby claim foreign priority benefits under 35 U.S.C. 119(a)-(d) or (f), or 365(b) of any foreign application(s) for patent, inventor's or plant breeder's rights certificate(s), or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below, by checking the box, any foreign application for patent, inventor's or plant breeder's rights certificate(s), or any PCT international application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application Number(s)	Country	Foreign Filing Date (MM/DD/YYYY)	Priority Not Claimed	Certified Copy Attached? Yes	No
			☐	☐	☐
			☐	☐	☐
			☐	☐	☐
			☐	☐	☐

☐ Additional foreign application numbers are listed on a supplemental priority data sheet PTO/SB/02B attached hereto.

PTO/SB/01 (06-03)
Approved for use through 07/31/2003, OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

DECLARATION — Utility or Design Patent Application

Direct all correspondence to: [✓] Customer Number: 23644 OR [ ] Correspondence address below

Name BARNES & THORNBURG		
Address PO BOX 2786		
City CHICAGO	State ILLINOIS	ZIP 60690-2786
Country UNITED STATES OF AMERICA	Telephone (312) 357 1313	Fax (312) 759 5646

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. 1001 and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

NAME OF SOLE OR FIRST INVENTOR: [ ] A petition has been filed for this unsigned inventor

Given Name (first and middle [if any]) ROBERT		Family Name or Surname FRISKNEY	
Inventor's Signature			Date
Residence: City HARLOW	State ESSEX	Country UNITED KINGDOM	Citizenship UNITED KINGDOM
Mailing Address 78 ALBERT GARDENS			
City HARLOW	State ESSEX	ZIP CM17 9QG	Country UNITED KINGDOM

NAME OF SECOND INVENTOR: [ ] A petition has been filed for this unsigned inventor

Given Name (first and middle [if any]) NIGEL		Family Name or Surname BRAGG	
Inventor's Signature			Date
Residence: City WESTON COLVILLE	State CAMBRIDGESHIRE	Country UNITED KINGDOM	Citizenship UNITED KINGDOM
Mailing Address HOMEWARDS CHAPEL ROAD			
City WESTON COLVILLE	State CAMBRIDGESHIRE	ZIP CB1 5NX	Country UNITED KINGDOM

[✓] Additional inventors or a legal representative are being named on the 1 supplemental sheet(s) PTO/SB/02A or 02LR attached hereto.

PTO/SB/02A (08-03)
Approved for use through 08/31/2003. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

DECLARATION

ADDITIONAL INVENTOR(S)
Supplemental Sheet

Page 1 of 1

Name of Additional Joint Inventor, if any: ☐ A petition has been filed for this unsigned inventor

Given Name (first and middle (if any))	Family Name or Surname
SIMON	PARRY

Inventor's Signature | Date

Residence: City	State	Country	Citizenship
HARLOW	ESSEX	UNITED KINGDOM	UNITED KINGDOM

Mailing Address: 2 KINGSTON COTTAGES, MATCHING GREEN

Mailing Address: 2 KINGSTON COTTAGES, MATCHING GREEN

City	State	Zip	Country
HARLOW	ESSEX	CM17 0PS	UNITED KINGDOM

Name of Additional Joint Inventor, if any: ☐ A petition has been filed for this unsigned inventor

Given Name (first and middle (if any))	Family Name or Surname
PETER	ASHWOOD-SMITH

Inventor's Signature: [signature] | Date: April 5th /2004

Residence: City	State	Country	Citizenship
HULL	QUEBEC	CANADA	CANADA

Mailing Address: 20 DES GENEVRIERS

Mailing Address: 20 DES GENEVRIERS

City	State	Zip	Country
HULL	QUEBEC	J9A 2V8	CANADA

Name of Additional Joint Inventor, if any: ☐ A petition has been filed for this unsigned inventor

Given Name (first and middle (if any))	Family Name or Surname
DAVID I	ALLAN

Inventor's Signature: [signature] | Date: April 5th /2004

Residence: City	State	Country	Citizenship
OTTAWA	ON	CANADA	CAN

Mailing Address: 852 FOREST ST.

Mailing Address: 852 FOREST ST.

City	State	Zip	Country
OTTAWA	ONTARIO	K2A 3P9	CANADA

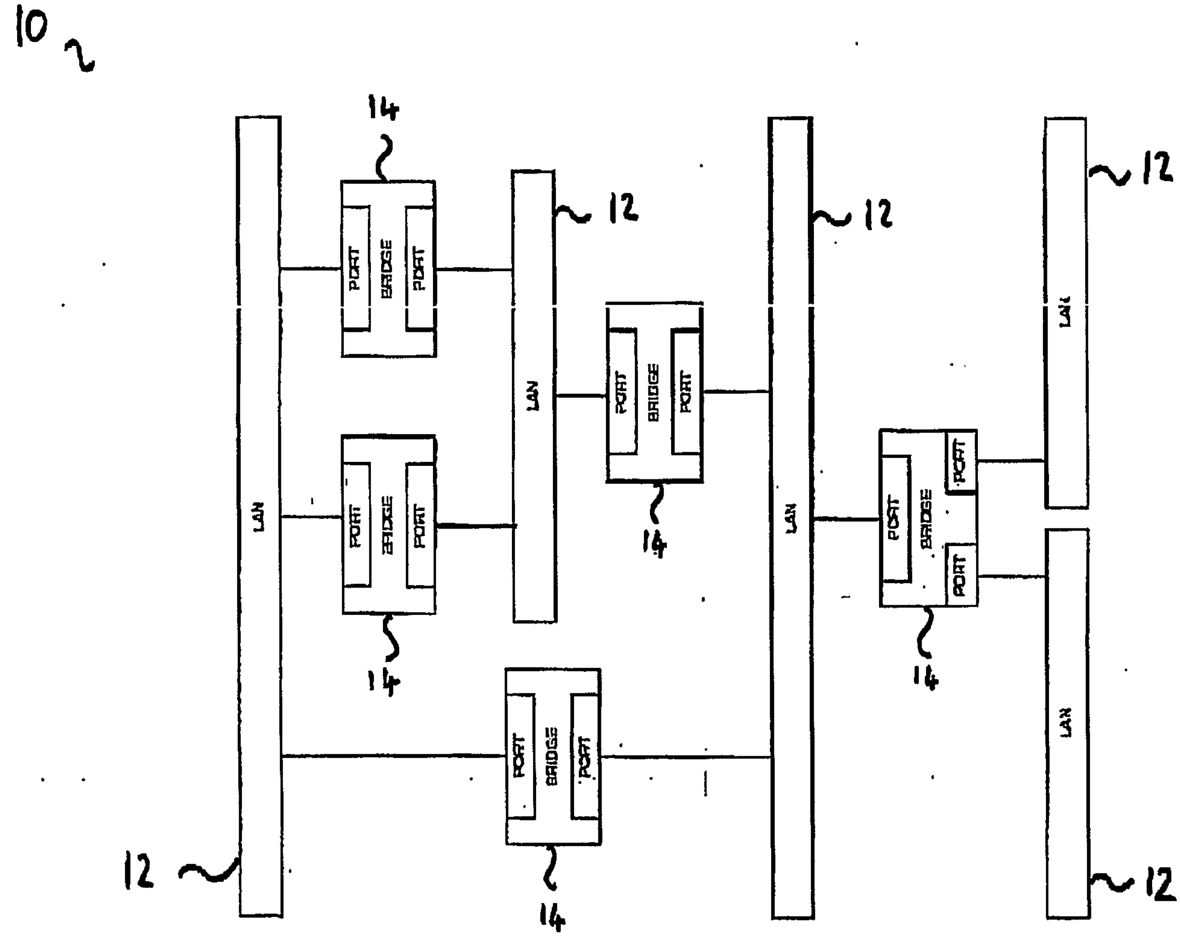

Prior Art

Fig 1

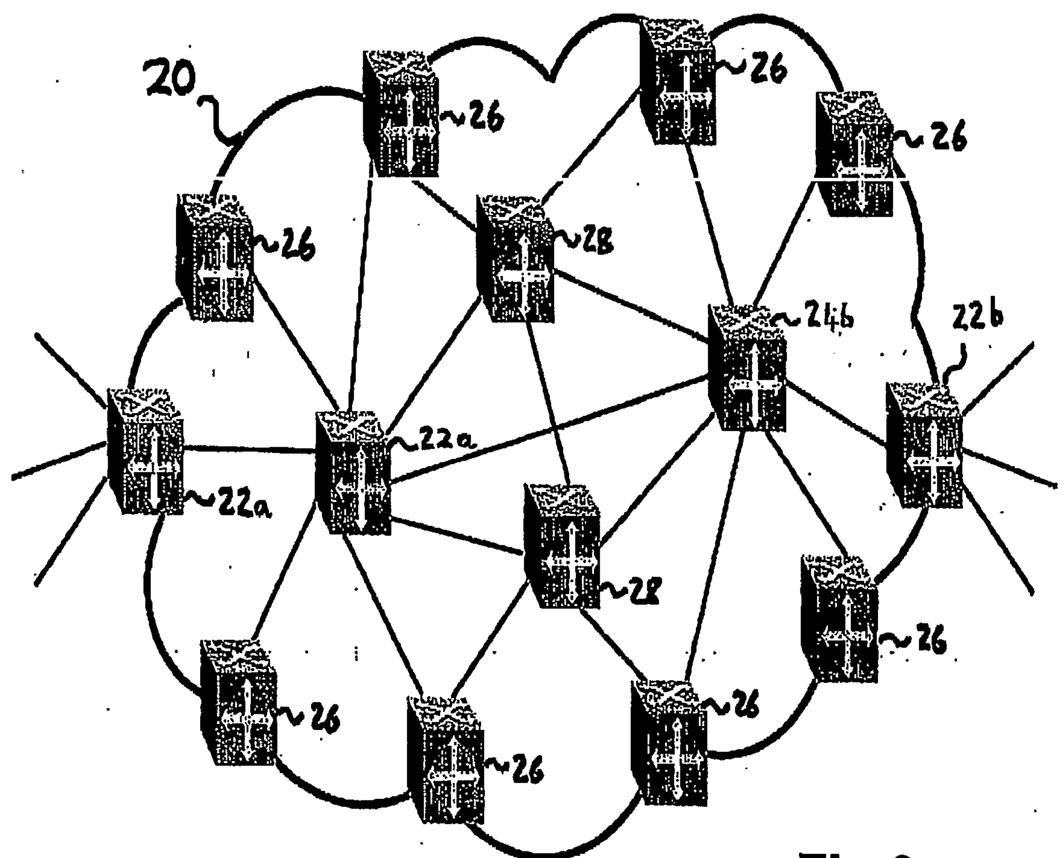

Fig 2

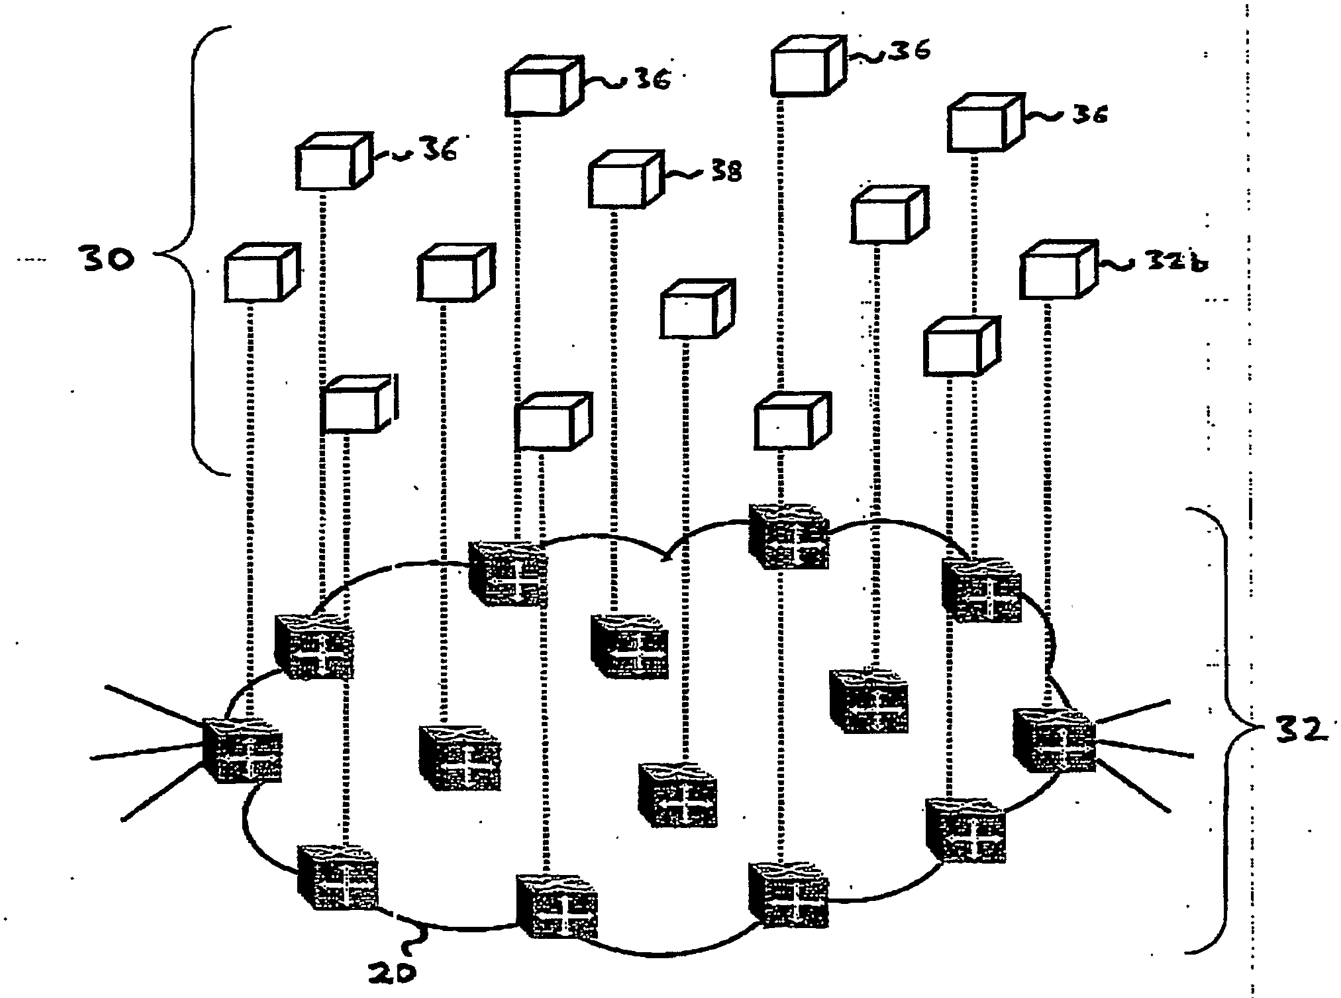

Fig 3

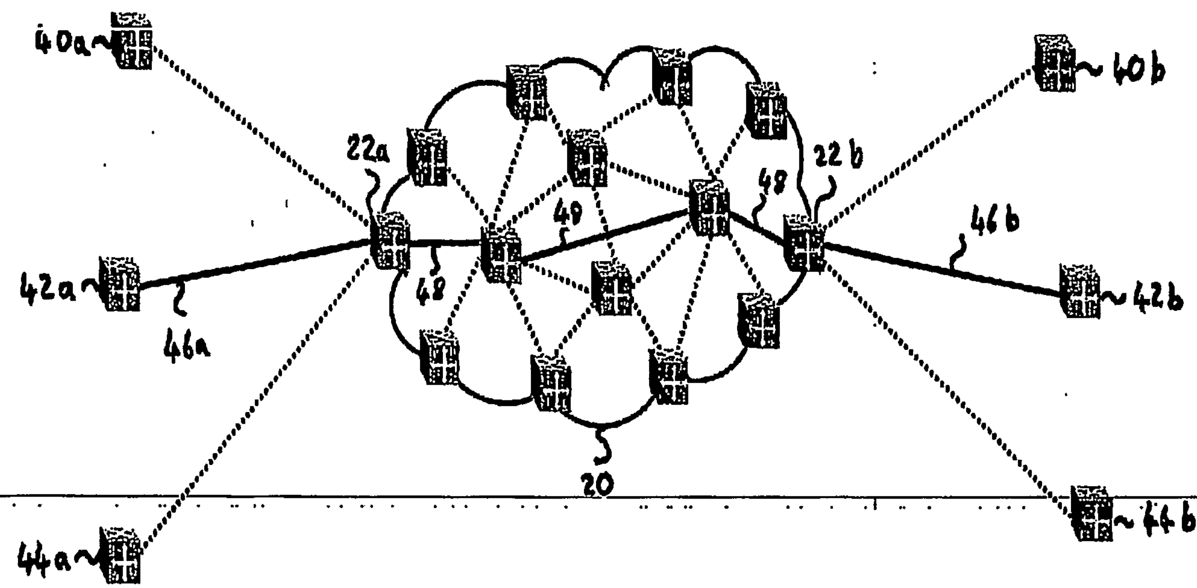

Fig 4

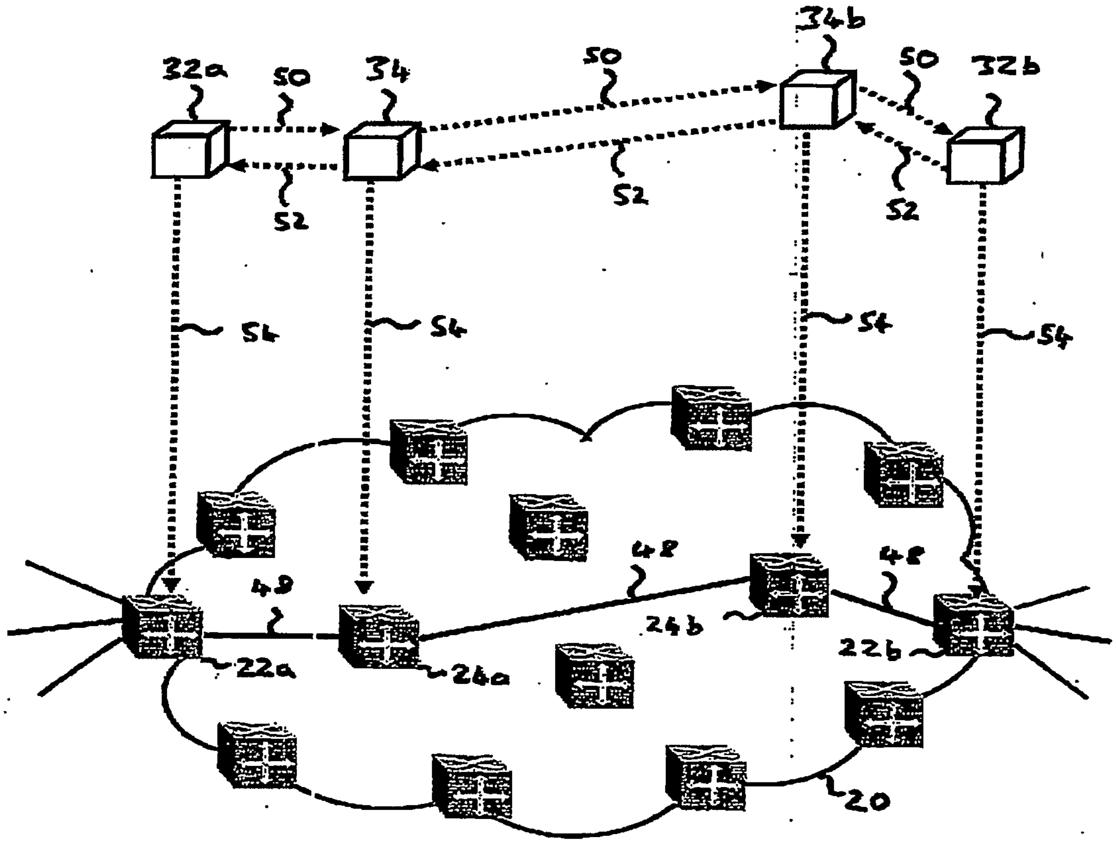

32a
50
34
50
34b
50
32b
52
52
52
54
54
54
54
48
48
48
24b
22b
22a
24a
20

Fig 5

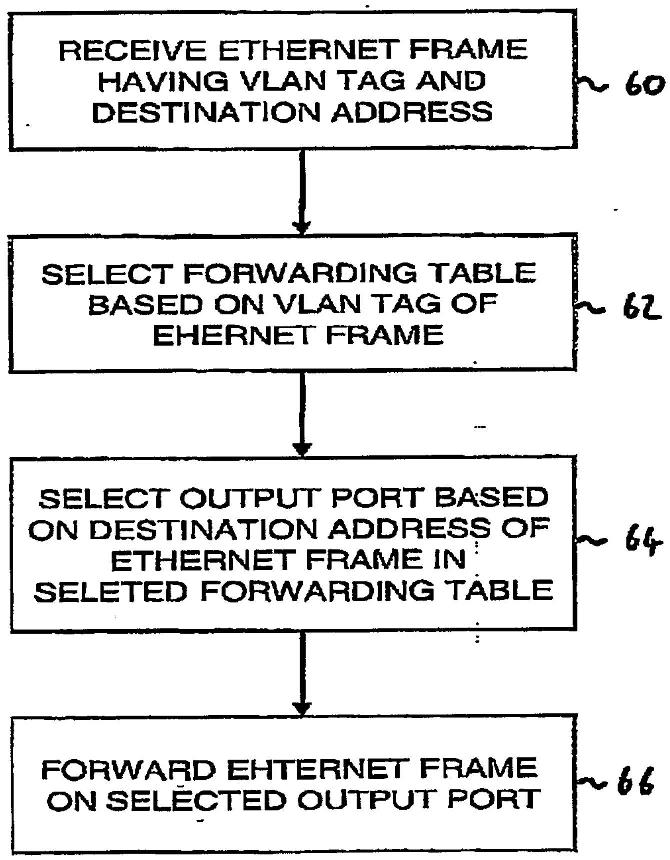
RECEIVE ETHERNET FRAME HAVING VLAN TAG AND DESTINATION ADDRESS
60
SELECT FORWARDING TABLE BASED ON VLAN TAG OF EHERNET FRAME
62
SELECT OUTPUT PORT BASED ON DESTINATION ADDRESS OF ETHERNET FRAME IN SELETED FORWARDING TABLE
64
FORWARD EHTERNET FRAME ON SELECTED OUTPUT PORT
66

Fig 6

Fig 7

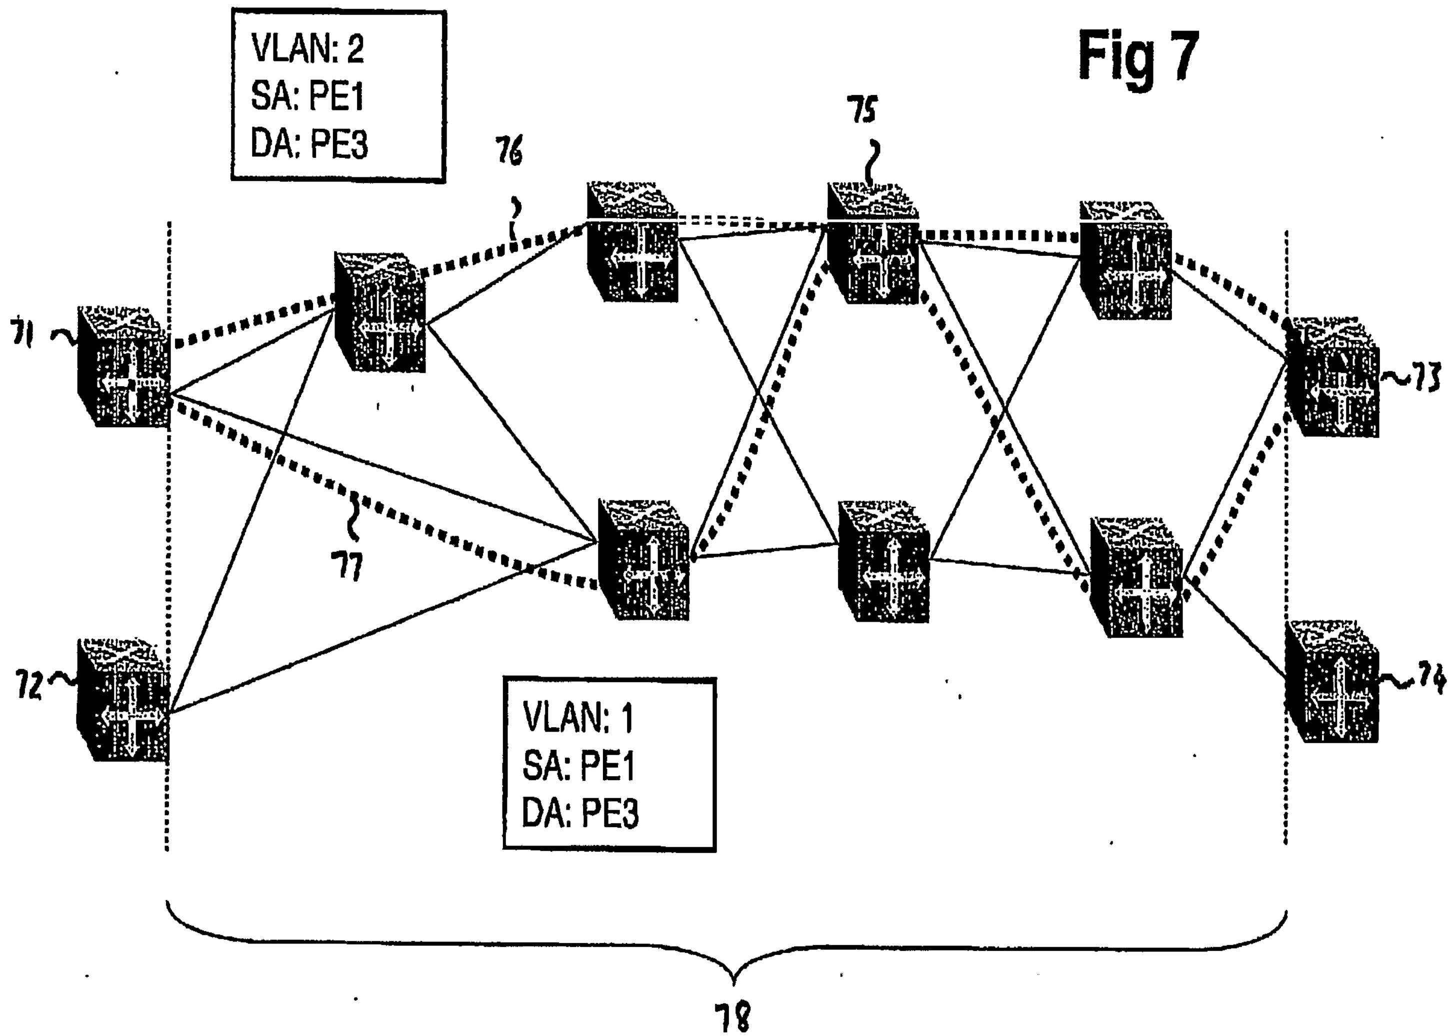

Fig 8

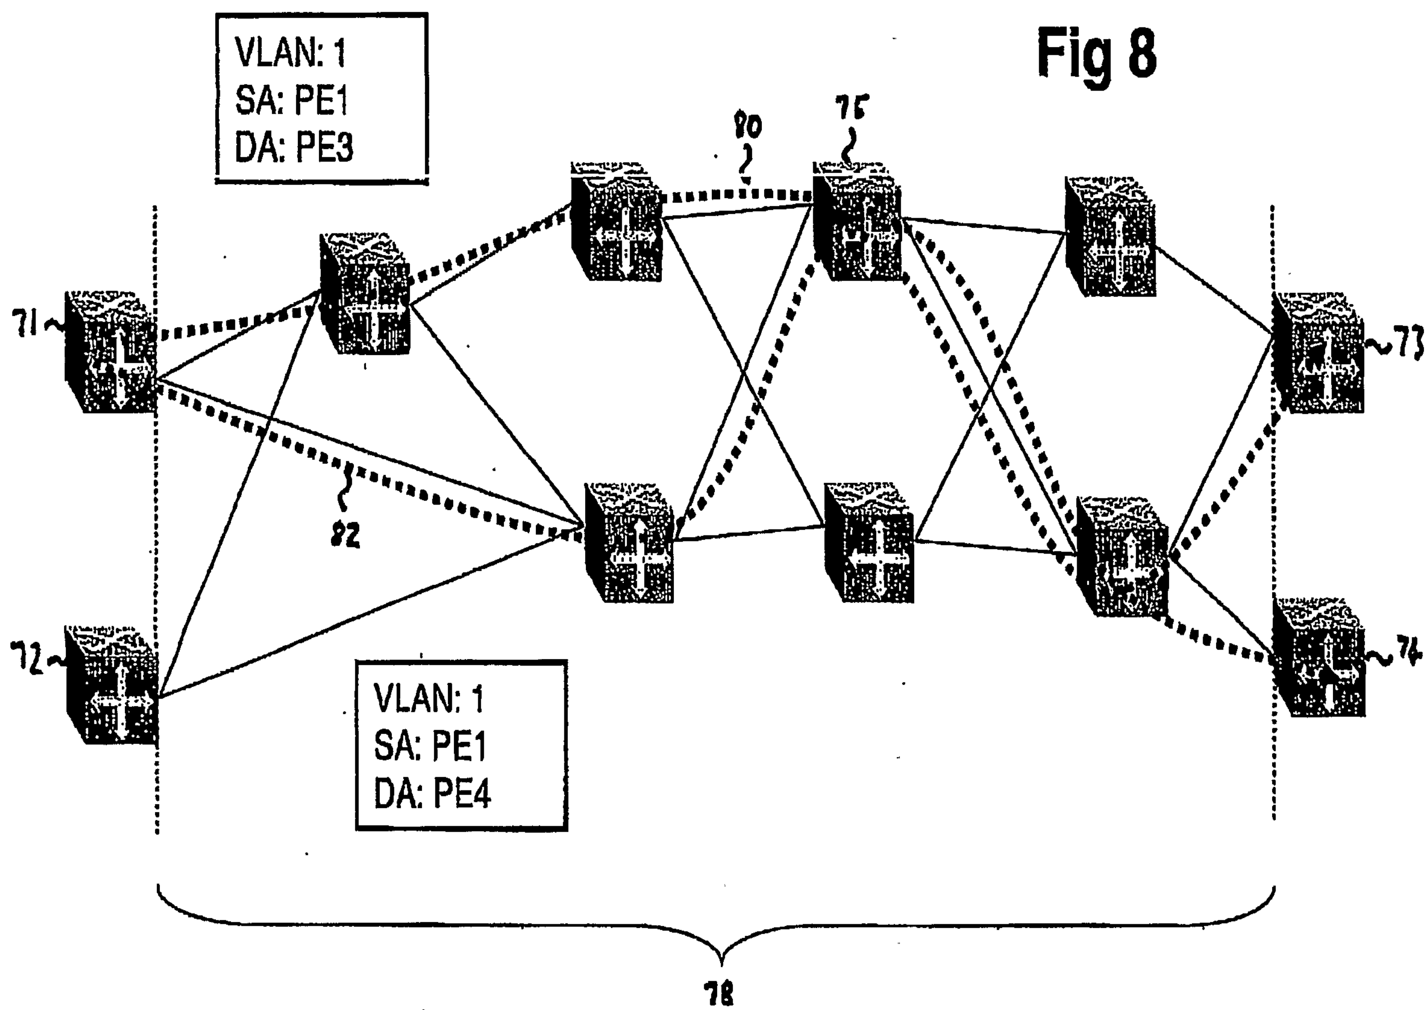

Fig 9

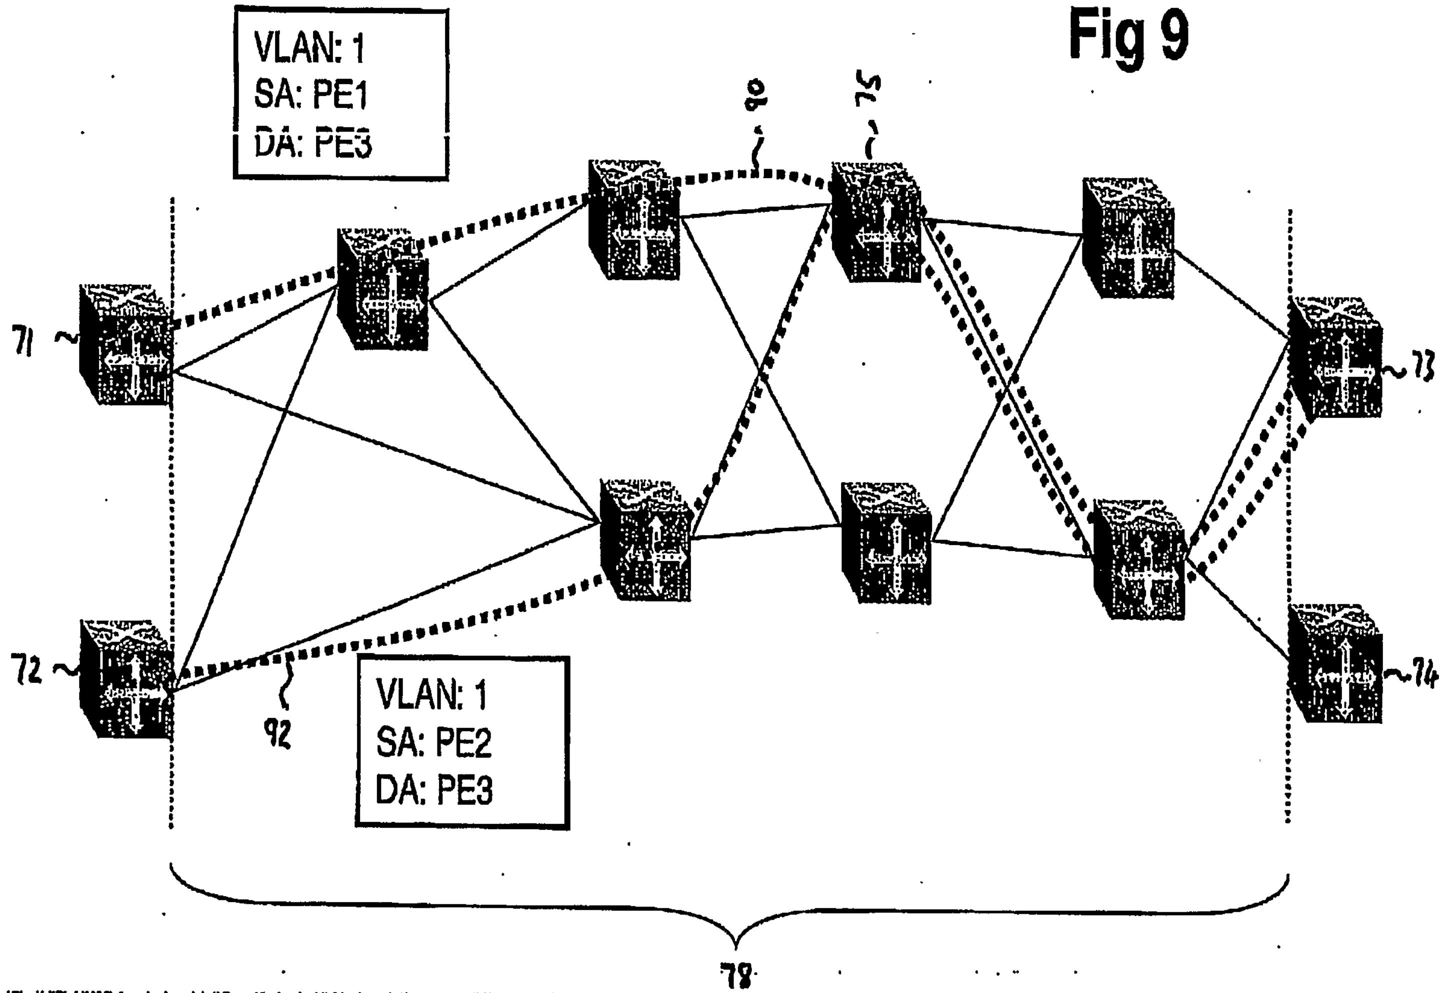

Fig 10

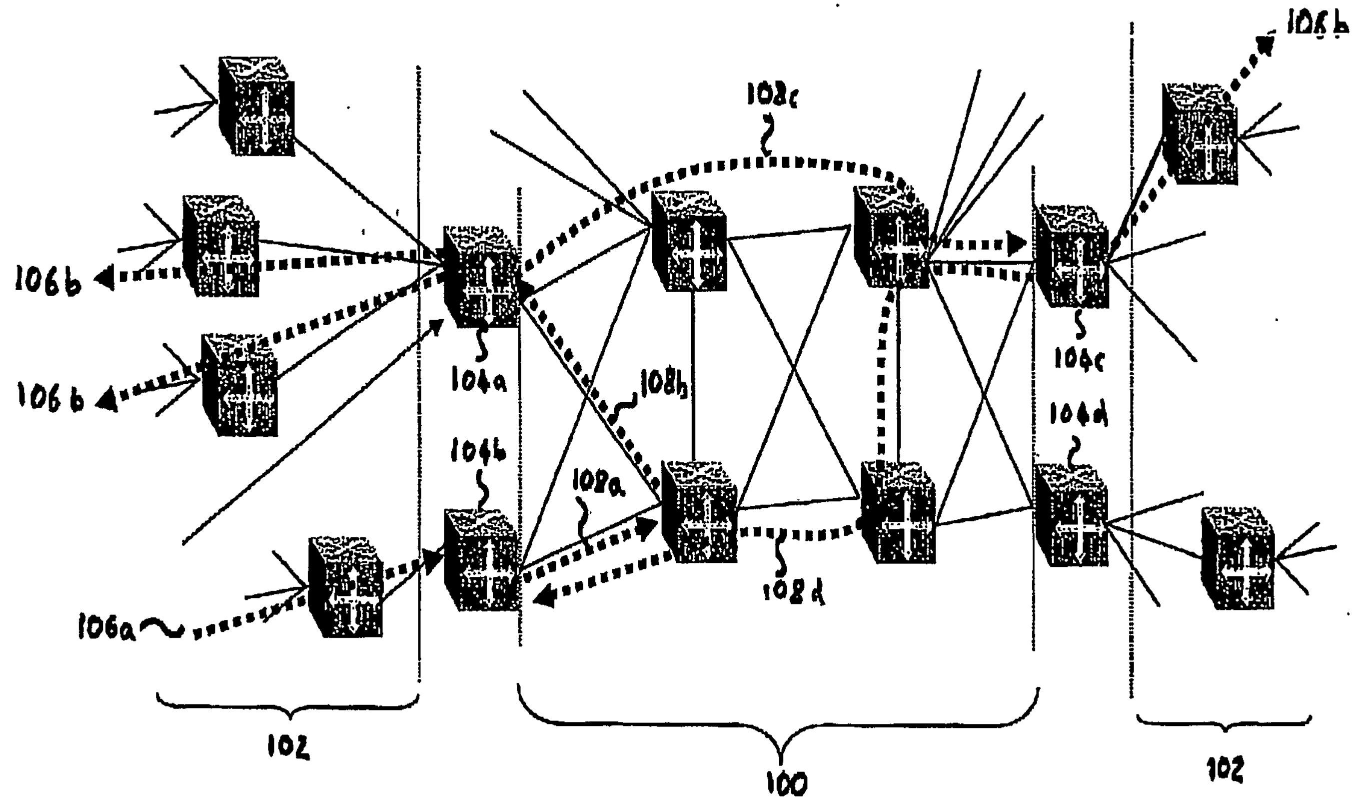
106b
106b
106a
104a
104b
108h
108a
108c
108d
104c
104d
106b
102
100
102

Fig 11

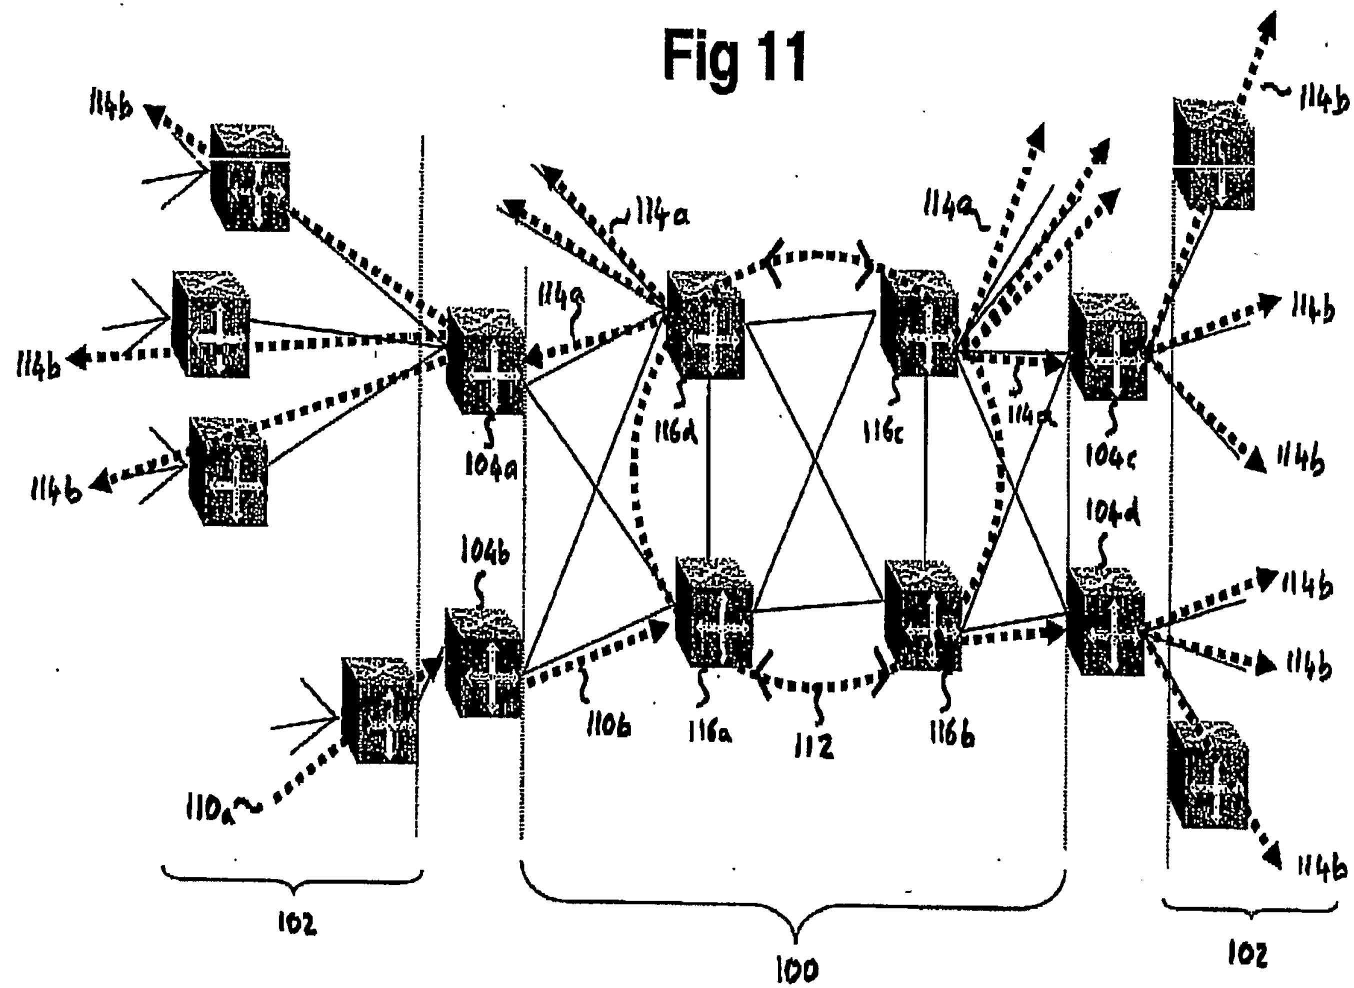
114b
114b
114b
114a
114a
114a
104a
104b
110a
110b
116d
116a
116c
116b
112
114a
114a
104c
104d
114b
114b
114b
114b
114b
114b
102
100
102